

THE AMENDMENTS

In The Claims

1. (Currently Amended) A system for communicating proprietary control information over one or more backplane connections, interconnecting two or more entities comprising a network entity without functioning as a user interface, comprising:

first logic for storing proprietary control information, comprising control information recognized by the network entity but not generally recognized by other network entities, within a layer of a packet above the physical layer; and

second logic for communicating the packet, including the proprietary control information, over the one or more of the backplane connections;

wherein the proprietary control information as stored in the packet either replaces or appears in the packet to ~~one or more other network entities~~ any third party devices that may happen to gain access as at least a portion of one or more standard packet fields;

wherein the first logic derives at least a portion of the control information from a packet header, and deletes the packet header prior to communication of the packet over the one or more backplane connections, and wherein a third logic re-creates at least a portion of the packet header from the control information after communication of the packet over the one or more backplane connections, wherein the derivations, deletions and re-creations of the packet header mitigate bandwidth limitations caused by the backplane connections and improve system throughput.

2. (Original) The system of claim 1 wherein the control information is stored in at least a portion of one or more fields inserted into the packet by the first logic.

3. (Original) The system of claim 1 wherein the first logic overwrites at least a portion of one or more pre-existing fields in the packet with the control information.

4. (Original) The system of claim 1 wherein the two or more entities comprise a switch, and the control information is proprietary to the switch.

5. (Original) The system of claim 4 wherein the switch has ingress and egress ports.
6. (Original) The system of claim 5 wherein the proprietary control information comprises an identifier of an ingress port of the switch at which the packet was received over a network.
7. (Original) The system of claim 5 wherein the proprietary control information comprises an identifier of an egress port of the switch at which the packet will or is expected to be transmitted over a network.
8. (Original) The system of claim 5 wherein the proprietary control information comprises an indicator of whether or not one or more predetermined fields were present in the packet upon receipt thereof at the switch.
9. (Original) The system of claim 8 wherein the one or more predetermined fields comprise a VLAN.
10. (Original) The system of claim 1 wherein the control information is stored in layer two or higher of the packet according to the OSI reference model.
11. (Original) The system of claim 10 wherein the control information is stored in layer two of the packet according to the OSI reference model.
12. (Original) The system of claim 11 wherein the control information is stored in the MAC sub-layer of the packet.
13. (Original) The system of claim 12 wherein the control information overwrites at least a portion of a VLAN stored in the MAC sub-layer of the packet.

14. (Original) The system of claim 12 wherein the control information overwrites at least a portion of source or destination addresses stored in the MAC sub-layer of the packet.

15. (Original) The system of claim 9 wherein the VLAN comprises op code and tag portions, and the first logic overwrites the op code portion of the VLAN with the control information.

16. (Original) The system of claim 15 wherein the control information comprises an identifier of the VLAN op code overwritten by the control information.

17. (Original) The system of claim 9 wherein the VLAN is the outer VLAN of a plurality of nested VLANs.

18. (Original) The system of claim 1 wherein the control information comprises quality of service information for the packet.

19. (Currently Amended) The system of claim 18 wherein the quality of service information comprises an identifier of a queue for buffering the packet and priority for the queue is selected based upon the quality of service information of the packet.

20. (Original) The system of claim 1 wherein the control information comprises an indicator that the packet is a candidate for dropping.

21. (Currently Amended) The system of claim [[3]] 1 wherein by dropping a portion of the packet, the control information is added to the packet and the packet is communicated in-band over the one or more backplane connections without requiring additional clock cycles.

22. (Canceled)

23. (Canceled)

24. (Currently Amended) A method of communicating proprietary control information over one or more backplane connections, interconnecting two or more entities comprising a network entity without functioning as a user interface, comprising:

storing proprietary control information, comprising control information recognized by the network entity but not generally recognized by other network entities, in a layer of a packet above the physical layer; and

communicating the packet, including the proprietary control information, over the one or more of the backplane connections,

wherein the proprietary control information as stored in the packet either replaces or appears in the packet to one or more ~~other network entities~~ any third party devices that may happen to gain access as at least a portion of one or more standard packet fields;

wherein at least a portion of the control information is derived from a packet header, and is deleted from the packet header prior to communication of the packet over the one or more backplane connections, and wherein at least a portion of the packet header from the control information is re-created after communication of the packet over the one or more backplane connections, wherein the derivations, deletions and re-creations of the packet header mitigate bandwidth limitations caused by the backplane connections and improve system throughput.

25. (Original) The method of claim 24 further comprising storing the control information in at least a portion of one or more fields inserted into the packet to accommodate the control information.

26. (Original) The method of claim 24 further comprising overwriting at least a portion of one or more pre-existing fields in the packet with the control information.

27. (Original) The method of claim 24 wherein the two or more entities comprise a switch, and the control information is proprietary to the switch.

28. (Original) The method of claim 27 wherein the switch has ingress and egress ports.

29. (Original) The method of claim 28 wherein the proprietary control information comprises an identifier of an ingress port of the switch at which the packet was received over a network.

30. (Original) The method of claim 28 wherein the proprietary control information comprises an identifier of an egress port of the switch at which the packet will or is expected to be transmitted over a network.

31. (Original) The method of claim 28 wherein the proprietary control information comprises an indicator of whether or not one or more predetermined fields were present in the packet upon receipt thereof at the switch.

32. (Original) The method of claim 31 wherein the one or more predetermined fields comprise a VLAN.

33. (Original) The method of claim 29 wherein the proprietary control information comprises an indicator of a state of the ingress port of the switch at which the packet was received.

34. (Original) The method of claim 24 wherein the control information is stored in layer two or higher of the packet according to the OSI reference model.

35. (Original) The method of claim 34 wherein the control information is stored in layer two of the packet according to the OSI reference model.

36. (Original) The method of claim 35 wherein the control information is stored in the MAC sub-layer of the packet.

37. (Original) The method of claim 36 wherein the control information overwrites at least a portion of one or more fields stored in the MAC sub-layer of the packet.

38. (Original) The method of claim 37 wherein the one or more fields comprise a VLAN.

39. (Previously presented) The method of claim 37 wherein the one or more fields comprise source or destinations addresses.

40. (Original) The method of claim 38 wherein the VLAN comprises op code and tag portions, and the control information overwrites the op code portion of the VLAN.

41. (Original) The method of claim 40 wherein the control information comprises an identifier of the VLAN op code overwritten by the control information.

42. (Original) The method of claim 38 wherein the VLAN comprises the outer VLAN of a plurality of nested VLANs.

43. (Original) The method of claim 24 wherein the control information comprises quality of service information for the packet:

44. (Currently Amended) The method of claim 43 wherein the quality of service information comprises an identifier of a queue for buffering the packet and priority for the queue is selected based upon the quality of service information of the packet.

45. (Original) The method of claim 24 wherein the control information comprises an indicator that the packet is a candidate for dropping.

46. (Currently Amended) The method of claim 24 wherein by dropping a portion of the packet, the control information is added to the packet and the packet is communicated in-band over the one or more backplane connections without requiring additional clock cycles.

47. (Canceled)

48. (Canceled)

49. (Original) The system of claim 5 wherein the switch is a first switch, and the proprietary control information comprises an identifier of an ingress port of a second switch coupled to the first switch at which the packet was received over a network.

50. (Original) The method of claim 28 wherein the switch is a first switch, and the proprietary control information comprises an identifier of an ingress port of a second switch coupled to the first switch at which the packet was received over a network.

51. (Currently amended) The system of claim 1 ~~further comprising~~ wherein the third logic for maintaining a mode bit having first and second states, wherein the first logic is configured to add one or more fields to the packet layer to accommodate the control information if the mode bit is in the first state, and overwrite at least a portion of one or more pre-existing fields in the packet layer with the control information if the mode bit is in the second state.

52. (Original) The method of claim 24 further comprising:
maintaining a mode bit having first and second states;
adding one or more fields to the packet layer to accommodate the control information if the mode bit is in the first state; and

overwriting at least a portion of one or more pre-existing fields in the packet layer with the control information if the mode bit is in the second state.

53. (Currently Amended) A system for performing load balancing over a plurality of backplane connections, interconnecting two or more entities comprising a network entity without functioning as a user interface, the system comprising:

first logic for receiving a packet at a first entity, mapping control information for the packet into one or more identifiers of at least one of the one or more backplane connections coupling the first entity to a second entity, the first and second entities comprising a network entity, wherein the mapping occurs through a data structure configured to achieve a desired load balancing of packets over the plurality of backplane connections; and

second logic for communicating the packet over the at least one of the one or more backplane connections identified by the one or more identifiers;

wherein quality of service information is used to identify a queue into which the packet is stored before transmission over the one or more backplane connections and priority for the queue is selected based upon the quality of service information of the packet.

54. (Original) The system of claim 53 wherein the two or more entities comprise a switch, and the control information is an identifier of an ingress port at which the packet was received over a network, or an egress port at which the packet will or is expected to be transmitted over a network.

55. (Previously Presented) The system of claim 54 wherein the first logic comprises a lookup table ("LUT") for maintaining an association between ingress or egress ports, and egress ports associated with the backplane connections, and the first logic maps a particular ingress or egress ports into one or more backplane-associated egress ports through an access to the LUT.

56. (Original) The system of claim 55 wherein the association is programmed into the LUT.

57. (Original) The system of claim 56 wherein the association is pre-determined to achieve a desired load balancing of packets over the plurality of backplane connections.

58. (Original) The system of claim 53 wherein the two or more entities are each ASICs.

59. (Currently Amended) A method of performing load balancing over a plurality of backplane connections, interconnecting two or more entities comprising a network entity without functioning as a user interface, the method comprising:

receiving the packet at a first entity coupled to a second entity through the one or more backplane connections, the first and second entities comprising a network entity;

mapping control information for a packet into one or more identifiers of at least one of the one or more backplane connections through a data structure configured to achieve a desired load balancing of packets over the one or more backplane connections; and

communicating the packet over the at least one of the one or more backplane connections identified by the one or more identifiers;

wherein quality of service information is used to identify a queue into which the packet is stored before transmission over the one or more backplane connections and priority for the queue is selected based upon the quality of service information of the packet.

60. (Original) The method of claim 59 wherein the two or more entities comprise a switch, and the control information comprises an identifier of an ingress port at which the packet was received over a network, or an egress port at which the packet will or is expected to be transmitted over a network.

61. (Previously Presented) The method of claim 60 further comprising using a lookup table ("LUT") to maintain an association between ingress or egress ports and egress ports

associated the backplane connections, and mapping an ingress or egress port into one or more of the backplane-associated egress ports through an access to the LUT.

62. (Original) The method of claim 61 further comprising programming the association into the LUT.

63. (Original) The method of claim 62 wherein the association is pre-determined to achieve a desired load balancing of packets over the plurality of backplane connections.

64. (Original) The method of claim 59 wherein the two or more entities are each ASICs.

65. (Currently Amended) A system for extending the number of ports of a switch in a network comprising:

a first switch coupled to a second switch and the first switch having a greater number n of ports than the number of ports m of the second switch;

first logic for storing in a layer of the packet above the physical layer an identifier of a port of the first switch;

second logic for communicating the packet between the first and second switches,

wherein the second switch appears to the network to have n ports rather than m ports.

66. (Original) The system of claim 65 wherein the port is an ingress port of the first switch at which the packet was received over a network, and the second logic communicates the packet from the first switch to the second switch.

67. (Original) The system of claim 65 wherein the port is an egress port of the first switch at which the packet will or is expected to be transmitted over a network, and the second logic communicates the packet from the second switch to the first switch.

68. (Original) The system of claim 65 wherein the port identifier is stored in layer two or higher of the packet according to the OSI reference model.

69. (Original) The system of claim 68 wherein the port identifier is stored in layer two of the packet according to the OSI reference model.

70. (Original) The system of claim 69 wherein the port identifier is stored in the MAC sub-layer of the packet.

71. (Original) The system of claim 70 wherein the port identifier is stored in the packet in the form of one or more standard fields.

72. (Original) The system of claim 70 wherein the port identifier is stored in the packet as a VLAN.

73. (Currently amended) A method of extending the number of ports of a switch in a network comprising:

providing a first switch coupled to a second switch and the first switch having a greater number n of ports than the number of ports m of the second switch;

storing in a layer of the packet above the physical layer an identifier of a port of the first switch; and

communicating the packet between the first and second switches,

wherein the second switch appears to the network to have n ports rather than m ports.

74. (Original) The method of claim 73 wherein the port is an ingress port of the first switch at which the packet was received over a network, and the packet is communicated from the first switch to the second switch.

75. (Original) The method of claim 73 wherein the port is an egress port of the first switch at which the packet will or is expected to be transmitted over a network, and the packet is communicated from the second switch to the first switch.

76. (Original) The method of claim 73 wherein the port identifier is stored in layer two or higher of the packet according to the OSI reference model.

77. (Original) The method of claim 76 wherein the port identifier is stored in layer two of the packet according to the OSI reference model.

78. (Original) The method of claim 77 wherein the port identifier is stored in the MAC sub-layer of the packet.

79. (Original) The method of claim 78 wherein the port identifier is stored in the packet in the form of one or more standard fields.

80. (Original) The method of claim 79 wherein the port identifier is stored in the packet in the form of a VLAN.

81. (Currently Amended) A system for communicating proprietary control information over one or more backplane connections, interconnecting two or more entities comprising a network entity without functioning as a user interface, comprising:

first means for mapping proprietary control information, comprising control information not generally recognized by network entities other than the network entity, for a packet into one or more identifiers of one or more of the plurality of backplane connections;

second means for storing the proprietary control information in a layer of the packet above the physical layer, wherein the proprietary control information as stored in the packet either replaces or appears in the packet to one or more other network entities as at least a portion of one or more standard packet fields; and

third means for communicating the packet, including the proprietary control information, over the identified one or more backplane connections;

wherein the first means derives at least a portion of the control information from a packet header, and deletes the packet header prior to communication of the packet over the one or more backplane connections, and wherein a second means re-creates at least a portion of the packet header from the control information after communication of the packet over the one or more backplane connections, wherein the derivations, deletions and re-creations of the packet header mitigate bandwidth limitations caused by the backplane connections and improve system throughput.

82. (Original) The system of claim 81 wherein the two or more entities comprise a switch, and the system further comprises means for extending the number of ports of the switch.

83. (Currently Amended) A method of communicating proprietary control information over one or more backplane connections, interconnecting two or more entities comprising a network entity without functioning as a user interface, comprising:

a step for mapping proprietary control information, comprising control information recognized by the network entity but not generally recognized by other network entities, for a packet into one or more identifiers of at least one of the one or more-backplane connections;

a step for storing the proprietary control information in a layer of the packet above the physical layer, wherein the proprietary control information as stored in the packet either replaces or appears in the packet to one or more other network entities as at least a portion of one or more standard packet fields; and

a step for communicating the packet, including the proprietary control information, over the at least one backplane connections identified by the one or more identifiers;

wherein at least a portion of the control information is derived from a packet header, and is deleted from the packet header prior to communication of the packet over the one or more backplane connections, and wherein at least a portion of the packet header from the control information is re-created after communication of the packet over the one or more backplane

connections, wherein the derivations, deletions and re-creations of the packet header mitigate bandwidth limitations caused by the backplane connections and improve system throughput.

84. (Original) The method of claim 83 wherein the two or more entities comprise a switch, and the method further comprises a step for extending the number of ports of the switch.

85. (New) The system of claim 1, further comprising one or more I/O blades interconnected with one or more management/switching blades;

wherein each I/O blade is coupled to a network and the I/O blade comprises one or more first MAC controllers which in turn are coupled to a first switch fabric, which in turn are each coupled to a first backplane MAC controller;

wherein the management/switching blades are coupled to the first backplane MAC controller, and the management/switching blade comprise one or more second backplane MAC controllers which are each coupled to a packet filtering and control engine, which in turn are each coupled to a second switch fabric, which in turn is coupled to a microcontroller.

86. (New) The system of claim 85 wherein the first switching fabric stores the packet in a queue associated with an egress port represented by a port tag indicator value for the packet, and wherein the queue is selected based upon the quality of service information for the packet.

87. (New) The method of claim 73, further comprising the first switch is a third party switch and the second switch is a proprietary switch;

wherein the process at a third party switch comprise the steps of:

receiving the packet at the third party switch;

determining if the packet is an ingress packet or an egress packet; and

processing the packet;

wherein if the packet is an ingress packet, the switch inserts a VLAN into the MAC sub-layer of the packet, and stores in the VLAN an identifier of the ingress port at which the packet was received by the third party switch, wherein the packet is transferred to the proprietary switch;

wherein if the packet is an egress packet, an identifier of the egress port is retrieved from the VLAN, a decision on deletion of the VLAN egress port identifier is

made and the packet is transmitted over the network from the designated egress port of the third party switch;

wherein the processing at a proprietary switch comprising the steps of:

- receiving the packet at the proprietary switch;
- determining if the packet is an ingress packet or an egress packet; and
- processing the packet;

wherein if the packet is an ingress packet, an identifier of the ingress port of the third party switch port based VLAN is obtained, the identifier is copied into a packet header, a decision on deletion of the port-based VLAN indicator is made, additional processing of the packet is preformed, and the packet is transferred to the MSM blade in the proprietary switch over one or more backplanes;

wherein if the packet is an egress packet, an identifier of the egress port is obtained from the packet header, a VLAN is inserted into the MAC sub-layer of the packet, and an identifier of the egress port of the third party switch is stored in the VLAN, and the packet is transmitted to the third party switch.

88. (New) A system of claim 1, further comprising a fourth logic that masks the control information such that the control information appears as a standard packet field to a third party.